An ADTRAN White Paper

Defining Broadband Speeds: an Analysis of Required Capacity in Network Access Architectures

Defining Broadband Speeds: An Analysis of Required Capacity in Network Access Architectures

Executive Summary

In Form 477, which is used by the FCC to collect data from service providers on broadband internet access, the FCC defines basic broadband internet connections as those that have an “information transfer rate” of at least 768 kbps in one direction (download or upload) [1]. Also in the recent rulemaking, the FCC inquires about whether to require “service providers to report actual measured speed, rather than the maximum possible speed, for each broadband connection.”

This paper supplements an earlier ADTRAN White Paper [2] and examines broadband speed requirements in terms of the total traffic forecast by Cisco’s Visual Networking Index [11], converting the total volume figures provided in the source into per-household metrics that can be used to estimate required capacity in smaller networks such as access networks. The metrics are extrapolated from the forecast trends to provide estimates of required capacity for the next six years.

Three different access network architectures – DSL, HFC, and WiMAX – are then presented and analyzed with respect to their per-household capacity in typical configurations. The analyses yield the following results:

- ADSL access networks designed to Carrier Serving Area guidelines will exceed the required capacities projected for year 2015. VDSL networks can provide additional capacity above that of ADSL.
- HFC access networks with 250 households per Optical Node and two downstream data channels meet current but not future capacity requirements. Providing an Optical Node per 125 households and increasing the number of downstream channels provides capacity the may become marginal around year 2015.
- Range-limited WiMAX such as those expected to serve rural areas are limited by the rate provided out to the cell edge. A range-limited cell with six sectors will support about 30 households at the required capacity for year 2015.

1 Introduction

In Form 477, which is used by the FCC to collect data from service providers on broadband internet access, the FCC defines basic broadband internet connections as those that have an “information transfer rate” of at least 768 kbps in one direction (download or upload) [1]. (Note that in the more recent rulemaking order (WC Docket No. 07-38 [FCC 08-89], released June 12, 2008), the terms “download speed” and “upload speed” are used in place of upload and download “information transfer rate.”) Also in the recent rulemaking, the FCC inquires about whether to require “service providers to report actual measured speed, rather than the maximum possible speed, for each broadband connection.”[1]

This paper examines broadband speed requirements in terms of the total traffic forecast by Cisco’s Visual Networking Index [11]. Trends for each of the applications sub-segments in the index are presented and the sub-segments are discussed with regard to the qualitative and quantitative requirements they place on networks. Different types of Internet video traffic and peer-to-peer traffic are given special consideration due to their trends, total volumes, and/or real time requirements.

The total traffic volumes provided in the index are then converted to estimates of per-household volumes, with scaling added to account for diurnal patterns. The average volumes are also scaled to account for non-uniform usage distributions and self-similar traffic distributions, resulting in capacity metrics that can be applied (with caution) in setting requirements for or planning future access networks.

Three types of access network architectures are then presented and analyzed for capacity against the capacity metrics. Typical DSL networks are shown to exceed the requirements extrapolated out to year 2015. Depending on the number of households served and channels dedicated to data, HFC networks will also meet the requirements projected to that year. WiMAX networks will meet the projected requirements if the deployments serve limited numbers of households – depending on cell range, from 20 to less than 10 households per sector.

2 Factors affecting traffic loading

Traffic loading and usage characteristics for High Speed Internet Access (HSIA) have increased in recent years and that increase is expected to continue. Some factors that affect traffic are examined below.

2.1 Applications

Much of the data on current and forecasted application traffic volumes in this section is taken from Cisco’s Visual Networking Index [11].

[1] This was left as a matter for further comment.

2.1.1 Internet Video

Internet video is defined as video content which is accessed over the Internet via a subscriber's HSIA service (as opposed to IPTV, which is sourced by a subscriber's service provider as a service separate from HSIA). While the best known example of Internet video is YouTube, there are many different sources, including broadcast and cable TV network web sites, social networking sites, movie delivery services, and educational web sites.

Internet video is widely considered to be the single largest factor in the growing requirement for bandwidth in broadband data services. Usage of the application has increased over 100% in the last two years [3] and is expected to triple by 2013 [4]. The growth of this application is changing traffic characteristics for HSIA in the following ways:

- It triggers a corresponding increase in raw volume. Current playout rates for Internet video range from approximately 300 kbps (YouTube standard definition) to approximately 2 Mbps ("HD" content from network web sites), for videos that may range from tens of seconds to hours in length. As true HD content at 5 to 6 Mbps becomes more widely available it will continue to drive volumes higher yet [12].
- The playout rate is almost always tied to a near-real time requirement for content delivery. Subject to the size of the receive buffer, viewing a video file in near-real time requires a data transfer rate at or above the playout rate, which must be sustained with little or no interruption for the duration of the video.
- The application is driving higher usage statistics. As people increasingly turn to Internet video instead of traditional sources for video entertainment, the percentage of subscribers who are actively using the service at a given time grows. For instance, from 2007 to 2008 the average video-to-PC consumption per household jumped by approximately 50% (per the analysis in section 3). Since the majority of that consumption was from YouTube [11], which did not initiate HD videos until December 2008, the majority of that increase can be attributed to a higher number of video streams rather than a higher average rate per stream.

Most sources classify Internet video differently from video content downloaded using peer-to-peer (P2P) applications for purposes of traffic analysis. The former is generally accessed via a client-server model, from sites such as YouTube, and watched in near-real time as it is being streamed (although some video content can be downloaded and stored for later viewing). P2P traffic, discussed in section 2.1.2, has different characteristics. Some applications such as Joost, discussed in section 2.1.3, combine aspects of both of the above categories.

2.1.1.1 Home networks and video to TV

The increasing proliferation of home networks is driving a corresponding increase in both the number of total subscribers on a given network and the usage statistics per subscriber household. The first and most obvious result of this proliferation is that more subscribers in a household can be online simultaneously on different computers. In addition, there is a growing market for devices that allow Internet video to be played out on a television rather than a computer monitor [6, 7, 8]. As this market matures, it will reinforce the growth of Internet video by making the viewing experience more familiar, regardless of the source (*e.g.*, families watching movies or TV episodes via Internet video rather than broadcast or cable).

In its Visual Networking Index report [11], Cisco breaks Internet video into the separate categories of video-to-PC and video-to-TV. Video-to-TV content is identified as "film and television content" as opposed to the user-generated content common on YouTube.

2.1.2 Peer-to-Peer (P2P)

P2P applications and protocols, and the resulting traffic loads, have been extensively analyzed [15, 16, 17]. Estimates of the amount of traffic generated by P2P applications in recent years have ranged as high as almost 80% of all consumer traffic [13]. One of the defining characteristics of P2P traffic is its symmetry – for every peer receiving content over a P2P network, another peer must be sending. Even though newer P2P protocols such as BitTorrent get content from multiple peers to reduce the peak upload burden on any one host, the nature of the application dictates that upload and download traffic is balanced over the population of the network. This can strain networks that were designed on the premise of client-server applications and asymmetric traffic loads.

A second characteristic of P2P traffic is that the difference between peak and average daily load levels tends to be less than for other applications [13]. Since most P2P applications deal with non-real time traffic, some users presumably schedule P2P transfers for non-peak traffic periods so as to not interfere with their interactive applications.

While the raw volume of P2P traffic continues to grow, its percentage share is steadily shrinking. Estimated P2P traffic for 2008 was 35% of North American consumer traffic, down from 61% in 2006 (and down further still from the 80% estimated in 2003 [13]). At least part of this trend results from the growing dominance of non-P2P video as a percentage of Internet traffic.

2.1.3 P2P video services (Joost)

Relatively new services like Joost [18, 19] combine some of the challenging characteristics of both streaming video and P2P applications. Joost enables subscribers to download video, including feature-length TV programs and movies, for near-real time viewing (as well as storage and later viewing) using a P2P protocol. As with BitTorrent and other P2P protocols, each file is transferred in "chunks" from multiple peers.

This class of application combines the near-real time requirement of streaming video with the symmetric traffic loading of P2P applications. To the degree that it is adopted, it will change access network requirements for both capacity and symmetry.

2.1.4 Video communications

While video communications (two-party video calls or video conferencing) is not yet widely adopted, it may finally be about to emerge as a significant application due to three enabling factors:

- Widespread broadband access,
- Widely available, inexpensive and easily installed webcams (frequently integrated in new laptops), and
- Free, widely available video communications features added on to VoIP and instant messaging applications.

While Cisco does not project video communications to take off before 2012, once growth does occur it will drive significant requirements for both symmetric and real time traffic volume.

2.1.5 Other applications

Other applications are noted below along with their expected effect on traffic demand in the access network.

- Traditional web browsing, email and file transfer applications will continue to represent a significant percentage of traffic volume. While speed is certainly a factor in the performance of these applications (especially file transfer), it is frequently less important that other factors such as latency [20].
- VoIP has widespread usage but, due to its low bit rate, it drives a small percentage of overall demand. Its main impact on access networks is that its performance is very dependent on congestion. Even momentary congestion will cause noticeable loss in voice quality.
- Gaming, like VoIP, has low bit rate requirements but stringent real time requirements. Also like VoIP, it can suffer noticeable loss of quality due to momentary congestion.

The monthly estimates and forecasts for North America Internet traffic by sub-segment for the years 2006 through 2012 are taken from different tables in [11] and compiled in Table 1.

Table 1 – North America Internet traffic by application class

By Sub-Segment (PB per month)	2006	2007	2008	2009	2010	2011	2012	CAGR 2007–2012
Web, email, data	152	209	280	365	478	620	799	31%
P2P	370	396	441	516	587	647	686	12%
Gaming	14	17	21	26	32	38	45	21%
Video communications	3	4	5	6	8	11	13	27%
VoIP	4	6	8	10	12	13	14	18%
Internet video to PC	59	186	317	406	506	635	771	33%
Internet video to TV	4	77	177	338	553	765	968	66%
Totals	606	895	1249	1667	2176	2729	3296	30%

2.2 Diurnal patterns

As is well documented [21, 13, 22], traffic volume exhibits a diurnal pattern reflecting user activity cycles. While business activity peaks during normal weekday office hours, consumer activity peaks during evening hours, with much less variation between weekdays and weekends. A diurnal pattern with similar peak times of day applies to different categories of traffic, although P2P traffic exhibits smaller excursions from the mean than interactive application classes [13]. Based on the levels of variation in [13], the overall mean traffic averaged over minutes during peak usage times can be expected to be at lease 1.5 times the long term average measured over days (see section 3).

2.3 Non-uniform usage

A common characteristic of Internet traffic is non-uniform demand across the subscriber population. As in many populations, a minority of subscribers consumes a majority of the measured traffic load. One study [13] indicated that 2.9% of subscribers (in a pool of over 100,000) accounted for over 40% of the traffic on the network, and that the top 20% of subscribers accounted for slightly over 80% of the traffic.

The concentration of demand in a small percentage of the subscriber population significantly increases the expected variance in demand in access networks, where the subscriber pool is smaller than in aggregation or core networks. In a network providing access to 100 subscribers, the usage characteristics of less than 20 households can be expected to dominate the traffic load.

2.4 Self-similar traffic

Another well documented characteristic of Internet traffic is self-similarity [23, 24]. One of the effects of self-similarity is fluctuations in momentary traffic load that exceed those predicted by models such as Poisson arrivals. Mori *et al.* [25] measured the skewness and marginal distributions of Internet traffic on a number of network links. The results show positively skewed distributions with momentary loads (summed and measured at 100 ms intervals) exceeding twice the mean rate in all measured traces.

2.5 Changing characteristics of Internet traffic

One final characteristic of Internet traffic is that it is easy to get into trouble trying to predict the future [14]. While some characteristics of the Internet are relatively invariant, such as diurnal patterns and traffic self-similarity, other characteristics can change almost literally overnight. New technologies, protocols and applications can reach critical mass and have a dramatic impact on traffic parameters.

The traffic volume forecast from Cisco [11] is based on extrapolation of current trends and best information available today, and they state that it "is considered conservative by many analysts." That forecast shows that video traffic in North America increased eightfold from 2006 to 2008, and that it is expected to increase by another 3.5x between 2008 and 2012. As one example, if video communications reaches critical mass earlier than expected we could see a significant change in both overall volume and symmetry of traffic relative to the forecast.

3 Analysis of traffic loading

The traffic volume values provided in [11] are monthly totals for consumer Internet traffic in North America. In this section we relate those figures to usage on a per-household basis in the belief that the resulting figures may provide some guidance for scaling shared capacity in access networks.

As a quick check on the source data, the total monthly volume reported by the Minnesota Internet Traffic Studies [26] for the US was from 1,200 to 1,800 Petabytes. This is in line with Cisco's estimate of 1250 Petabytes for monthly consumer traffic in North America in 2008.

The following discussion is based on 2008 figures. Based on US population estimates and the most recent census figures for persons per household [27], there were approximately 117 million households in the US. Approximately 55% of US adults had broadband Internet access and another 10% had dial-up access [28]. Assuming that the Pew survey did not include more than one person per household, we can infer a high correlation between personal and household Internet access since access is normally provided on a per-household basis. So, approximately 65 million households had broadband connections, which should account for the vast majority of consumer traffic (with broadband having 5.5 times the number of dialup connections at over 10 times the speed and longer average session times, we can safely assume that dialup volume was relatively minor).

From Table 1, total volume in 2008 was about 1250 Petabytes per month. Spreading 1250 Petabytes per month across 65 million households gives us a long term average volume (measured over days) of approximately 60 kbps per household.

Table 2 – Internet traffic long term average traffic per household

Estimated broadband adoption	2006	2007	2008	2009	2010	2011	2012
No. of households (million)	115	116	117	119	120	121	122
Broadband adoption rate	42%	47%	55%	59%	62%	65%	68%
No. of broadband households (million)	48.4	54.7	64.6	69.9	74.3	78.6	83.0
Traffic by Sub-Segment (kbps per household)							
Web, email, data	9.70	11.8	13.4	16.1	19.9	24.3	29.7
P2P	23.6	22.4	21.1	22.8	24.4	25.4	25.5
Gaming	0.893	0.960	1.00	1.15	1.33	1.49	1.67
Video communications	0.191	0.226	0.239	0.265	0.332	0.432	0.483
VoIP	0.255	0.339	0.382	0.441	0.499	0.510	0.520
Internet video to PC	3.76	10.5	15.2	17.9	21.0	24.9	28.7
Internet video to TV	0.255	4.35	8.46	14.9	23.0	30.0	36.0
Totals	38.7	50.5	59.7	73.6	90	107	123

Table 2 shows the above analysis applied to the data in Table 1. The broadband adoption rate for years 2006-2008 is from [28]. The rate for 2009 is 4% higher than 2008 based on 3% gains reported over 9 months, and the rate for 2010-2012 is extrapolated at a linear increase of 3% per year based on survey results in [28] indicating that the rate of broadband growth may be slowing.

The above data includes both upstream and downstream traffic averaged over an extended period. We use some additional data to estimate upstream and downstream volumes during peak daily periods. [13] provides approximate maximum/minimum load ratios for the daily traffic patterns for traffic from different applications, and notes that the peaks occur at about the same times across all applications. Modeling the diurnal excursions from the mean as approximately symmetric[2], the corresponding average loads during peak traffic hours can be estimated as

$$P = M\left(1 + \frac{r-1}{r+1}\right), \tag{1}$$

where: P = the average load during peak periods,
M = the long term average load, and
r = the diurnal max/min traffic ratio.

The same study provides upstream vs. downstream traffic ratios for traffic from different applications. These values are incorporated for the year 2008 data in Table 3.

Table 3 – Traffic during peak hours, 2008

Application class	Web	P2P	Video to PC	Video to TV	Other (1)	Total
Long term mean M (kbps)	13.4	21.1	15.2	8.46	1.63	59.7
Diurnal max/min r (2)	5	2	5	2	4	
Mean volume P (peak time)	22.3	28.1	25.3	11.3	2.6	89.5
Down/up ratio (3)	8	1	8	8	1	
% downstream	89%	50%	89%	89%	50%	
Downstream (peak time)	19.8	14.1	22.4	10.0	1.30	67.7
Upstream (peak time)	2.48	14.1	2.81	1.25	1.30	21.9

Notes on Table 3:

1. The Other class includes the gaming, video communications, and VoIP sub-segments.
2. [13] states that the maximum to minimum diurnal load ratio is about 2 for P2P traffic and about 5 for Web browsing traffic. For this analysis, the Web browsing ratio is applied to interactive categories and the P2P ratio is applied to categories in which files can be scheduled for off-peak download. Video to PC, which consists primarily of shorter clips at lower bit rates, is places in the interactive category. Video to TV, which includes feature length films at high playout rates, is placed in the off-peak category. While all the sub-segments in the Other class are interactive, VoIP and video calling may be somewhat more distributed in time so the ratio applied is reduced slightly.

[2] This assumption of symmetry probably results in underestimation of the peak period averages. The diurnal patterns for consumer traffic in [13] look approximately symmetric, but those in [22] look like they exhibit positive skewness, which would make the peak period volumes somewhat higher than those calculated here.

3. Downstream/upstream ratios in [13] are approximately 8 for client/server applications that primarily download data, and approximately 1 for symmetric applications. For this analysis, all video traffic is assumed to follow the client/server model. Increased adoption of P2P video (e.g., Joost) could push upstream rates higher.

An interesting point can be inferred from Table 3. The value shown for the average downstream video-to-PC traffic during peak hours in 2008 is 22.4 kbps. During most of that period, the most popular source for streaming video was YouTube and the average playout rate for the streaming content was about 300 kbps. Depending on factors such as how many households had multiple viewers accessing different content, the ratio between playout rates and average traffic indicates that during peak hours approximately 7% of households with broadband access were viewing streaming video at any given time.

Table 3 shows how values for traffic loading during peak usage times, as they would be measured over reasonably short time frames of several minutes, are derived. Estimates for the same parameters are provided in Table 4 for the years covered by the current Cisco forecast.

Table 4 – Average traffic scaled per household during peak usage hours

Direction	2006	2007	2008	2009	2010	2011	2012	CAGR 2007-2012
Down (kbps per household)	37	54	68	85	110	130	150	22%
Up (kbps per household)	19	21	22	25	29	33	35	11%

At this point, we need to step back and list the accumulated caveats regarding the above numbers.

- Despite the scaling, the average values shown in Table 4 are obviously not intended to be applied to individual households. Subject to the caveats below, they are scaling factors applicable (with the addition of judicious margins) to large numbers of households in a given population.
- As noted in section 2.3, actual traffic loads in a network are dominated by a small percentage of users. In a portion of an access network spanning 100 subscribers, half of the total load may be generated by a handful of users. Depending on design parameters, it might be necessary to add margin of 100% or more to cover variation in the user population.
- From section 2.4, we know that short-term traffic fluctuations in the aggregated traffic regularly exceed twice the average rate. In order to prevent frequent congestion in the access network from limiting performance of VoIP and other real time applications sensitive to congestion, planned capacity should include an additional 100% margin to cover these fluctuations.
- The extrapolation of future broadband adoption in this analysis may deviate significantly from the assumptions made by Cisco when generating their total volume forecasts. Differences in those extrapolations could have a significant impact on the per-household CAGR values.

- Finally, we need to remember that the forecast data used as the source of this analysis may not reflect actual future trends, considering the volatile history of the Internet. Rapid adoption of new applications, protocols or technologies could render the forecast obsolete, even in the limited three year future span that this forecast covers.
- Compared to the dire warnings listed above, the following items seem trivial, but they are included for completeness. The analysis used forecast numbers for North America (including Canada) against population figures for the United States only, which inflated the per-household figures by an estimated 10%. Partly offsetting that, the forecast values do not include signaling traffic, acknowledged in [11] to add about 3% to overall volumes.

The per-household averages in Table 4 must be scaled further, both to account for non-uniform demand in the user population and to account for momentary peaks in demand. As noted in section 2.3, user demand is highly non-uniform and may follow a Pareto or similar distribution. While such distributions can show self-similar characteristics, user demand is also bounded in actual networks, either by physical network elements or by a rate cap imposed by the access network provider. Preliminary simulations indicate that rate caps tend to limit the self-similarity in the user distribution so that the resulting variance tends toward being inversely proportional to the size of the user population, with required margins ranging from about 1.5 (for larger access networks on the order of 1000 users) to as high as about 4 (for smaller access networks on the order of 50 users). Since the actual user distributions aren't well known, however, we will use a simpler rule of thumb and scale the per-household averages by a factor of two to account for non-uniform demands. Note that this factor may be low relative to the actual variation in smaller networks.

Section 2.4 discusses the ratio between average and momentary peaks in demand. Due to the self-similar nature of network traffic, this ratio (on the order of two to one, based on the data in [25]) should not change significantly over the range of population sizes encountered in access networks.

If we apply the scaling factors discussed above to the data in Table 4, we arrive at a current (2009) required capacity on the order of 350 kbps downstream and 100 kbps upstream per household for shared resources in the access network. Projecting into the future, required capacity in 2012 is on the order of 600 kbps downstream and 150 kbps upstream per household. Given the uncertain adoption schedule of known potential symmetric applications such as P2P video streaming and video communications (not to mention the increasing potential for unforeseen changes over time), the 11% CAGR applied to upstream loading seems particularly conservative. A doubling of per-household traffic every three years in both directions would be a slightly more conservative (and perhaps better) rule of thumb. Applying that rule would give us the following required capacities for shared bandwidth per household over time (Table 5).

Table 5 – Approximate required capacity/household for shared facilities in the access network

Direction	2009	2012	2015
Down (kbps per household)	350	700	1400
Up (kbps per household)	100	200	400

The values for year 2015 in the above table should be considered tentative, since they extrapolate the CAGR three years beyond the forecast numbers provided by Cisco. Given the expense and time associated with deploying broadband infrastructure, however, a projected requirement that looks only three years into the future would result in deployments that are obsolete soon after their introduction.

4 Analysis of access architectures against required capacities

In each section below, the access network architecture is described and then per-household capacities are compared to the values in Table 5.

4.1 Digital Subscriber Loop

4.1.1 Architecture

A DSL access network is an example of an access architecture where the last-mile channel (in this case, the DSL link) is dedicated to a single subscriber. The DSL access architecture comprises two or more stages between the Internet Service Provider's (ISP) point of presence and the subscriber, as shown in Figure 1. They are:

1. The backhaul network[3] between the ISP and the DSL Access Multiplexer (DSLAM), located either in the central office or in the loop plant. When the DSLAM is located in the Central Office (CO), the network-facing connection of today's DSLAMs is generally a high speed data network operated by the access provider, with data rates at or above the Gigabit per second range.
2. The subscriber loop. The loop provides a dedicated connection to each subscriber from the DSLAM.

[3] In the architecture descriptions, and in the analyses which follow, we will neglect core network equipment such as core and edge routers since these points are independent of congestion in the access network architecture.

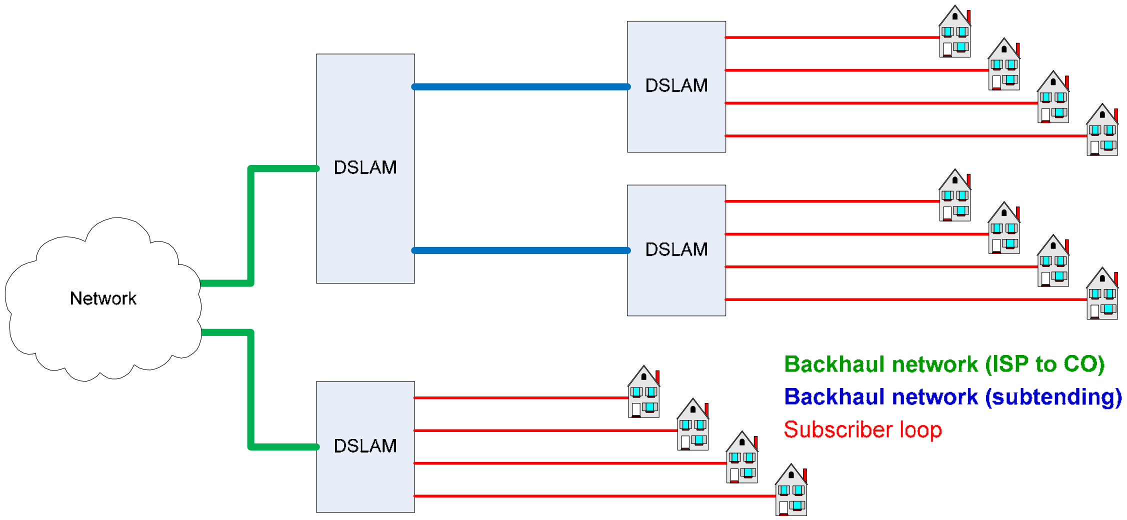

Figure 1 – DSL access network architecture

A typical CO-based DSLAM is a modular unit that may be populated with different types of access cards. Connections to the data network usually include multiple Gigabit or higher rate links. On the subscriber side, the DSLAM may support 500 or more DSL links, either directly or through subtended DSLAMs as described below.

In many DSL networks, remote DSLAMs are deployed in the loop plant. This decreases the length of the loop between the subscriber and the DSLAM, which in turn enables higher data rates. These DSLAMs usually serve from 24 to 384 subscribers in a Distribution Area (DA).

If the subscriber is served by a subtended DSLAM, there will be a connection between the CO-based DSLAM and the subtended DSLAM. Many subtended DSLAMs are fed over fiber links at Gigabit rates. Smaller DSLAMs may be fed by multiple copper loops using loop bundling.

The network-facing DSLAM connections described above are shared resources, over which data from multiple subscribers share bandwidth. These are typically high speed resources, operating in the Gigabit per second and above range.

Each subscriber loop connection is a point-to-point link between the DSLAM and a single subscriber. All traffic transmitted across that loop is dedicated to the subscriber served by the loop. With currently-available commercial technology, achievable rates on the longest loops of a Carrier Serving Area (12,000 ft) are approximately 6 Mbps for download and approximately 1 Mbps for upload, with much higher rates attainable on shorter loops. When loops are served by a remote DSLAM dedicated to a single distribution area, the maximum loop length is typically less than 6000 feet, supporting download data rates of 15-25 Mbps per subscriber.

4.2 Hybrid Fiber-Coaxial

An HFC cable access network is an example of an access architecture where the last-mile channel is a shared resource. Data transmission on an HFC network uses the DOCSIS protocol. An HFC access network typically comprises three connections between the network's point of connection to the Internet and the subscriber (Figure 2). They are:

1. The backhaul network between the ISP and the Cable Modem Termination System (CMTS) at the cable network head end or a hub site. As in a DSL network, this is generally a high speed network with data rates at or above the Gigabit per second range.
2. The fiber connection between the CMTS and an Optical Node. The Optical Node performs a conversion between optical and electrical signals for download traffic, and the inverse conversion from electrical to optical signals for upload traffic.
3. The coaxial network from the Optical Node to the pool of subscribers served by that node. Each coaxial network can serve up to 2000 subscribers in a tree and branch topology. In a modern network which includes data service, the coaxial network is typically sized on the order of 250 to 500 subscribers.

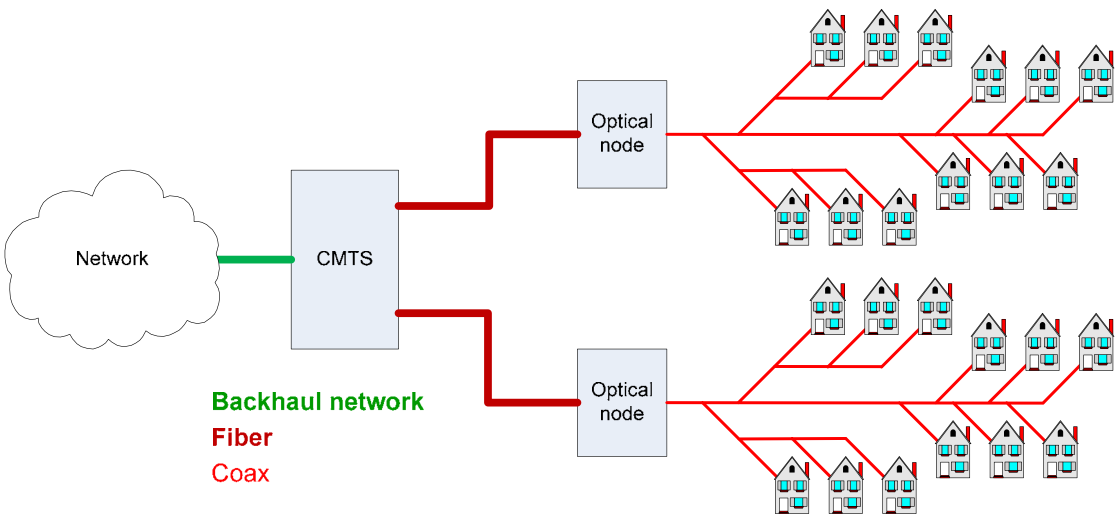

Figure 2 – HFC access network architecture

The signal transmission format is the same in the fiber and coaxial portions of the HFC network. In the download direction, data is modulated as an RF signal in one or more channel bands and multiplexed with analog and digital video in their own channel bands. The download spectrum includes 52 MHz to 760 MHz (some systems extend this range to 860 or 1000 MHz) and is divided into 6 MHz channels. The majority of these 6 MHz channels are used for delivery of television signals, while several of the channels are used for data transmission by the CMTS and cable modems (CMs). Multiple subscribers share the same 6 MHz channel for data transmission.[4] Within this channel, the download data RF signal is broadcast to all subscriber CMs, each of which decodes only the data intended for it.

[4] DOCSIS 3.0 allows several of the 6 MHz channels to be inverse multiplexed (or bonded) into a 'super' channel, which allows higher peak rates per subscriber. However, if the number of 6 MHz channels allocated to serve a given number of subscribers does not change, the average rate per subscriber is unaffected by the channel bonding.

Upload data, like download, is RF modulated and multiplexed into fixed channels. The upload spectrum includes 5 MHz to 42 MHz and, depending on the version of the DOCSIS in use, may be divided into channels of 6 MHz or smaller increments.[5] Unlike the download path, the upload path must merge data from many different sources onto the shared transmission channel. This is generally accomplished using Time Division Multiple Access (TDMA), although some versions of DOCSIS specify Synchronous Code Division Multiple Access (S-CDMA) as an option.

Under DOCSIS 2.0, usable shared data rates are up to 38 Mbps per download channel and up to 27 Mbps per 6 MHz upload channel. A typical residential deployment allocates one or two download channels to data [9]. Issues such as noise funneling and RF noise ingress tend to impose practical limits on both upload channel bandwidth and transmission density, resulting in a total shared upload capacity in current systems on the order of 35 Mbps for networks serving 250 subscribers [5]. DOCSIS 3.0 adds the capability to bond data from multiple channels together to increase peak rates, although it does not increase the per-channel rate in either direction.

In addition to the shared channel limits, the rate realized by the subscriber may be limited by the data rate that can be sustained by a single cable modem. This is much more likely to be a limit in the download than the upload direction.

4.3 Broadband Wireless Access

Broadband wireless access (BWA) is another example of shared last-mile access. With broadband wireless access, a number of subscribers share a wireless spectrum allocation, using a multiple access protocol to share the channel. While a number of BWA deployments have made use of either WiFi (IEEE 802.11b/g) or proprietary technologies, many wireless deployments going forward are based on WiMAX.[6] WiMAX specifies a set of profiles for wireless transmission based on IEEE 802.16e-2005[7] and related standards. While the IEEE standards define a large set of options, the profiles defined by WiMAX specify a subset of features that compliant systems must implement to ensure interoperability.

A WiMAX-based access network comprises at least two connections between the network's point of connection to the Internet and the subscriber (Figure 3). They are:

1. The backhaul network between the ISP and the WiMAX base stations. This network includes a high speed connection to the ISP and backhaul connections to the base stations. Some backhaul connections are wireless, using a point-to-point

[5] DOCSIS 3.0 adds an option to increase the high end of the upstream band from 42 to 85 MHz. When this option is used, the low end of the download band moves from 52 to 108 MHz. However, this requires a frequency band-plan change for all services, eliminating television channels 2-5.

[6] Carriers have also announced plans to use the next evolution of the 3GPP mobile wireless protocols known as LTE (Long Term Evolution) for broadband access. Although the specific analysis was based on WiMAX, the basic analysis presented here is extendable to most multiple access wireless protocols.

[7] While WiMAX implementations based on 802.16e-2005 are popularly known as "mobile WiMAX," and most deployments are targeting mobility, both the technology and the deployments serve fixed broadband purposes as well.

WiMAX connection. Others use fiber, DSL, or any of a number of other connection technologies.

2. The wireless network between the base station and the subscribers.

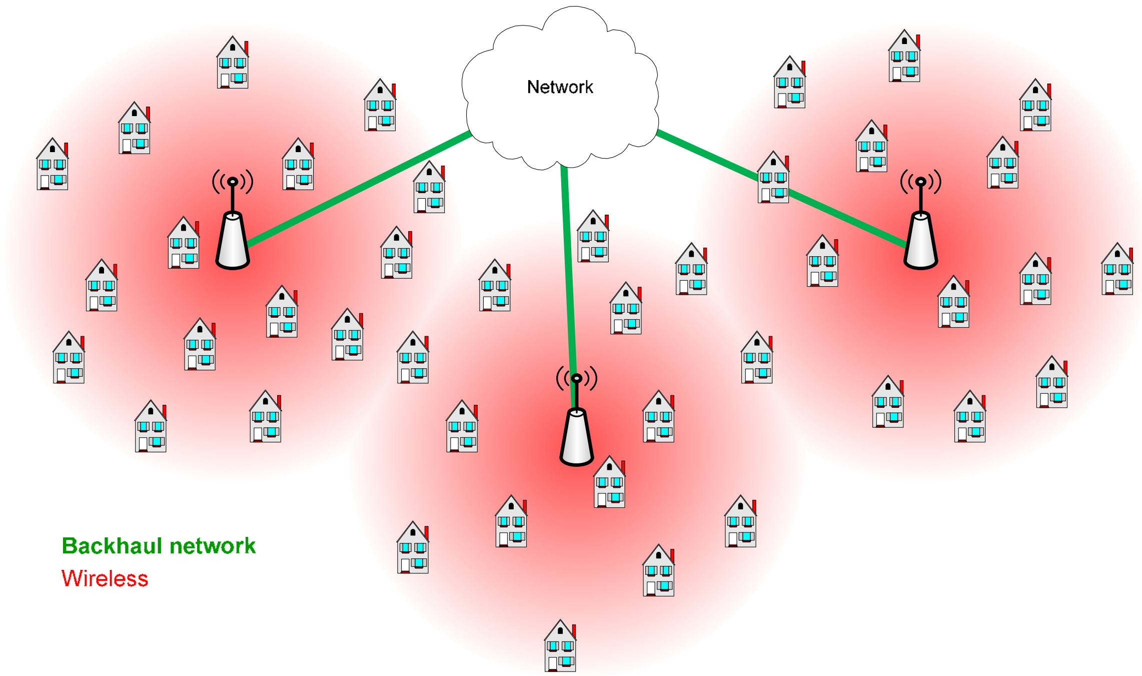

Figure 3 – Wireless access network architecture

The details of the shared channel vary by region. A notable deployment of WiMAX in the US is using the Broadband Radio Service (BRS) spectrum at 2.5-2.7 GHz. While multiple channel bandwidths are allowed in the 802.16e standard, the use of 5 MHz or 10 MHz channel bandwidths is common. WiMAX profiles for fixed broadband deployments allow either Time or Frequency Domain Duplexing (TDD or FDD). Current profiles for mobile deployments (which can also support fixed subscribers) specify only TDD, and even most fixed deployments use TDD. Since TDD uses the same channel for upload and download transmission and spends part of each time slot transmitting in each direction, the effective shared rate in either direction is reduced by the proportion of the time spent transmitting in the other direction and the guard time required while switching directions. While the theoretical maximum shared rate on a 10 MHz channel can approach 50 Mbps, the payload rate (split between upload and download) after subtracting PHY and MAC layer overhead is about 38 Mbps. Only subscribers closest to the base station will see that performance, with the shared rate decreasing with distance and obstructions between the transmitter and receiver. Note that because the channel is shared, the available bandwidth of the shared channel is effectively reduced to the average rate achievable by the active subscribers. This may be considerably lower than the peak rate available to subscribers near the base station.

WiMAX TDD transmission uses Orthogonal Frequency Division Multiple Access (OFDMA) organized in 5 msec radio frames. OFDMA allows upload and download data from different subscribers to be multiplexed in both the time and frequency domains. Resource allocation is defined on a per-frame basis using MAP fields, which have a variable length component that increases with the number of subscribers being scheduled [10].

WiMAX deployments may be either range limited or capacity limited. Deployments in rural areas are more likely to be range limited, with widely spaced base stations and relatively few subscribers per sector. In this scenario, the rate for a subscriber may be defined as much or more by the signal loss between the Customer Premise Equipment (CPE) and the base station as it is by how many subscribers are trying to share the channel.

Deployments in urban areas are more likely to be capacity limited, with closer spacing of base stations, multiple sectors per base station, and many subscribers per sector. In this scenario, a much higher percentage of CPEs can communicate with the base station at the highest shared rates, but subscribers must contend for upload and download bandwidth in much the same way as HFC-based subscribers.

As with HFC, in addition to the shared channel bandwidth limits, there may also be limits to the data rate that a single subscriber terminal can sustain.

5 Performance

Per-household rate capacities for the different types of access network architectures are derived in the sections below.

5.1 DSL performance

Three DSL network configurations are analyzed. The first represents a typical ADSL network designed to CSA parameters, with DSL loop rates configured to 6 Mbps download and 1 Mbps upload. Five 96-port remote DSLAMs, serving a total of 480 subscribers, are located in the loop plant and fed by 1 Gbps fiber uplinks to the CO-based DSLAM. The CO-based DSLAM supplies fiber to each subtended DSLAM and also directly serves an additional 20 subscribers located near the CO, bringing the total number of subscribers served to 500. The northbound network connection for the CO-based DSLAM is 1 Gbps.

The second configuration represents a VDSL deployment. The network topology is the same, but the DSL loop rates are now configured to 20 Mbps download and 4 Mbps upload.

The third configuration modifies the VDSL deployment by adding a second Gigabit link to the northbound CO interface, for a total capacity of 2 Gbps.

The relevant parameters and capacity for each network configuration are shown in Table 6. In each case except ADSL upstream, capacity per household is limited by the network backhaul link. Comparing the VDSL A and VDSL B cases shows that the capacity per household can be upgraded by adding to the backhaul network.

Table 6 – DSL parameters

Network	Direction	Connection	Resource type (# households)	Total capacity	Capacity per household
ADSL	Both	ISP to CO	Shared (500)	1 Gbps	2 Mbps
	Both	Subtending	Shared (96)	1 Gbps	10.4 Mbps
	Down	Subscriber loop	Dedicated	6 Mbps	6 Mbps
	Up	Subscriber loop	Dedicated	1 Mbps	1 Mbps
	Down	**Capacity per household**			**2 Mbps**
	Up	**Capacity per household**			**1 Mbps**
VDSL A	Both	ISP to CO	Shared (500)	1 Gbps	2 Mbps
	Both	Subtending	Shared (96)	1 Gbps	10.4 Mbps
	Down	Subscriber loop	Dedicated	20 Mbps	20 Mbps
	Up	Subscriber loop	Dedicated	4 Mbps	4 Mbps
	Down	**Capacity per household**			**2 Mbps**
	Up	**Capacity per household**			**2 Mbps**
VDSL B	Both	ISP to CO	Shared (500)	2 Gbps	4 Mbps
	Both	Subtending	Shared (96)	1 Gbps	10.4 Mbps
	Down	Subscriber loop	Dedicated	20 Mbps	20 Mbps
	Up	Subscriber loop	Dedicated	4 Mbps	4 Mbps
	Down	**Capacity per household**			**4 Mbps**
	Up	**Capacity per household**			**4 Mbps**

5.2 HFC performance

Two HFC network configurations are analyzed. The first network (HFC A) serves 250 customers off each Optical Node, with two RF channels allocated to download data. DOCSIS 3.0 channel bonding is enabled for a shared download rate of 76 Mbps. The shared upload rate is 35 Mbps.

The second network (HFC B) adds additional Optical Nodes so that each shared resource now serves only 125 subscribers. Two additional download channels are converted from video to data, which with channel bonding brings the total shared rate to 152 Mbps. The smaller split in the cable plant allows better performance in the upload direction, improving the shared rate to 52 Mbps.

In both network configurations, the network-to-CMTS connection allocates 1 Gbps per 500 subscribers. The upload bandwidth lost to the TDMA multiple access protocol is ignored for the purposes of this analysis, so these numbers may be slightly higher than actual rate available. The relevant parameters and capacities for each configuration are shown in Table 7.

Table 7 – HFC parameters

Network	Direction	Connection	Resource type (# households)	Total capacity	Capacity per household
HFC A	Both	Backhaul	Shared (500)	1 Gbps	2 Mbps
	Down	Fiber/coax	Shared (250)	76 Mbps	0.304 Mbps
	Up	Fiber/coax	Shared (250)	35 Mbps	0.140 Mbps
	Down	**Capacity per household**			**0.304 Mbps**
	Up	**Capacity per household**			**0.140 Mbps**
HFC B	Both	Backhaul	Shared (500)	1 Gbps	2 Mbps
	Down	Fiber/coax	Shared (125)	152 Mbps	1.22 Mbps
	Up	Fiber/coax	Shared (125)	52 Mbps	0.416 Mbps
	Down	**Capacity per household**			**1.22 Mbps**
	Up	**Capacity per household**			**0.416 Mbps**

5.3 Wireless performance

Analyzing expected performance for a WiMAX deployment is more complex than for either DSL or HFC, because there are more variables. Since the transmission medium is wireless, there is no physical infrastructure to define the number of subscribers served by a given base station or sector. Each access provider determines the design parameters for their network based on population density, expected take rate, licensed or unlicensed spectrum available, topology of the area to be served, and many other factors.

As noted before, some networks will be range limited while others will be capacity limited. In both types of networks, some CPEs will experience better signal path characteristics (and hence better overall rate performance) than others – this will be true to the largest degree in range limited networks, but because of signal degradation due to obstructions it is a significant factor in urban capacity limited networks as well.

Finally, shared bandwidth in each direction in a TDD network is dependent on the upload vs. download split of the traffic. This parameter is dynamic, changing on a frame-by frame basis.

A detailed analysis of WiMAX capacity would have to account for the above factors as well as the link budget and path loss in each direction, variable overhead dependent on the number of users, topology of the area to be served, variable population densities, the locations of heavy users relative to light users, and other factors. Instead, a very simplified capacity analysis is presented that assumes a fixed uplink/downlink symbol ratio (22 downstream symbols, 15 upstream symbols[8]) and overhead (11 symbols) in order to use the payload capacities tabulated in [29] for different constellations and code rates. Those capacities (for a 10 MHz channel with 2x2 MIMO) are reprinted in Table 8 along with the relative distances to which each modulation operates based on Free Space

[8] The ratio of uplink vs. downlink symbols may well be more evenly split (as shown) than would be predicted by the traffic ratio, since upstream reach at a given modulation rate may not be as far as the corresponding downstream reach.

Path Loss (FSPL). The propagation model assumes a flat, unobstructed coverage area for simplicity.

Table 8 – Capacities and relative working distances for WiMAX modulations

Region index *i*	Constellation	Code rate	Downlink rate (Mbps)	Uplink rate (Mbps)	Relative distance d_i
1	64 QAM	3/4 CTC	28.51	15.12	1.00
2	64 QAM	2/3 CTC	25.34	13.44	1.14
3	16 QAM	3/4 CTC	19.01	10.08	1.84
4	16 QAM	1/2 CTC	12.67	6.72	2.95
5	QPSK	3/4 CTC	9.5	5.04	3.85
6	QPSK	1/2 CTC	6.34	3.36	5.96

In the analysis of total capacity below, the frame mapper allocates bandwidth to all users regardless of distance to the base station. This method evenly distributes rate capacity throughout the sector and preserves per-user capacity out to the greatest distance from the base station.

Assuming a uniform population distribution across the geographic area, the number of households in the region served by a given modulation is proportional to the area of the annular ring (or the circle, for the innermost region) formed by the region. The proportion is the same regardless of the number of sectors into which the cell is divided.

$$N_i = N_T * \frac{d_i^2 - d_{i-1}^2}{d_T^2} \qquad (2)$$

For a fair traffic distribution (assuming uniform demand across the user population[9]), the total capacity Ca_i allocated to each region must be proportional to the number of users in that region.

$$Ca_i = Ca_T * \frac{N_i}{N_T} \qquad (3)$$

The percentage of capacity used in each region is its allocated capacity divided by the total capacity C_i for that region (where C_i equals the uplink or downlink rate deliverable to that region). That percentage is equal to the fraction of the WiMAX frame payload allocated to the users in that region. The sum of the percentage capacities over all regions should equal 100%.

$$\sum_i \frac{Ca_i}{C_i} = 1 \qquad (4)$$

Substituting the right side of Equation 3 for Ca_i in Equation 4 and rearranging terms solves for Ca_T, the total allocated capacity across the user population.

---

[9] As discussed in section 2.3, this assumption is known to be false. However, it is used here to enable a very basic model of wireless coverage across the sector.

$$Ca_T = 1 \Big/ \sum_i \frac{N_i}{C_i N_T} \tag{5}$$

Table 9 provides the total capacity available to a sector based on Equation 5. We see from the results that as the cell radius increases and less dense, more robust modulation schemes become necessary at the cell edge, the overall capacity is driven down rapidly by the increasing fraction of the TDD frame required to provide bandwidth to the outermost users.

Table 9 – Total allocated capacity across sector

# regions	Modulation at cell edge	Relative cell radius	Downlink capacity (Mbps)	Uplink capacity (Mbps)
1	64 QAM, 3/4 CTC	1.00	28.51	15.12
2	64 QAM, 2/3 CTC	1.14	27.73	14.71
3	16 QAM, 3/4 CTC	1.84	21.59	11.45
4	16 QAM, 1/2 CTC	2.95	15.10	8.01
5	QPSK, 3/4 CTC	3.85	12.15	6.45
6	QPSK, 1/2 CTC	5.96	7.92	4.20

Table 10 applies the required capacities from Table 5 for the year 2015 to the above results, to determine how many households could be supported by each sector. Downstream and upstream capacities need to be considered separately, since the breakpoints between regions are dependent on antenna gains and other factors and will in general be different for the downlink and uplink directions.

Table 10 – Number of households served across sector required capacity

# regions	Downlink		Uplink	
	Max users supported	Capacity per user (kbps)	Max users supported	Capacity per user (kbps)
1	20	1426	37	409
2	19	1460	36	409
3	15	1439	28	409
4	10	1510	20	400
5	8	1519	16	403
6	5	1584	10	420

Table 11 applies the above results to a range-limited network with six sectors per cell, which supports up to 30 households at year 2015 required capacities (ignoring factors such as frequency reuse and cell edge interference).

Table 11 – WiMAX network parameters

Network	Direction	Connection	Resource type (# households)	Total capacity	Capacity per household
WiMAX range limited (QPSK at cell edge, 6 sectors/cell)	Both	Backhaul	Shared (30)	1 Gbps	33.3 Mbps
	Down	Wireless	Shared (30)	47.5 Mbps	1.58 Mbps
	Up	Wireless	Shared (30)	25.2 Mbps	0.84 Mbps
	Down	**Capacity per household**			**1.58 Mbps**
	Up	**Capacity per household**			**0.84 Mbps**

6 Summary

Broadband speed requirements are examined in terms of the total traffic forecast by Cisco in their Visual Networking Index [11]. Trends for each of the applications sub-segments in the index are presented and the sub-segments are discussed with regard to the qualitative and quantitative requirements they place on networks. Different types of Internet video traffic and peer to peer traffic are given special consideration due to their trends, total volumes, and/or real time requirements.

The total traffic volumes provided in the index are then converted to estimates of per-household volumes, with scaling added to account for diurnal patterns. The resulting averages are then further scaled to provide margins for non-uniform usage distributions and self-similar traffic distributions. The resulting per-household capacity scalers are presented as figures that could be used (with suitable caution) in setting requirements for or planning future access network buildouts. The forecast traffic scalers are extrapolated to the year 2015 in the hopes that they will outlive the introduction phases of the networks currently being planned. The results of the analysis are repeated below.

Table 5 – Approximate required capacity/household for shared facilities in the access network

Direction	2009	2012	2015
Down (kbps per household)	350	700	1400
Up (kbps per household)	100	200	400

Three types of access network architectures are then presented and analyzed for capacity against the above values, with the following results:

- ADSL and VDSL access networks each serving 500 households are shown. In each case, the per-household capacities exceed the values required for the year 2015 with minimum capacities of 2 Mbps per household downstream and 1 Mbps per household upstream. In most cases (and in all of the VDSL cases), the per-household capacity is determined by the bandwidth in the backhaul network, rather than by the subscriber loop.
- HFC access networks with Optical Nodes serving 250 and 125 households are examined. In the 250 household case, the downlink and uplink capacities of 304

and 140 kbps per household respectively are marginal compared to 2009 requirements. The 125 household case, with downlink and uplink capacities of 1220 and 416 kbps per household, should meet traffic load requirements through approximately 2015.

- WiMAX networks were analyzed to determine total sector capacity and households supported per sector based on the modulation required to serve households furthest from the base station. The simplified analysis assumes uniform geographic distribution and demand, and uses the optimistic FSPL model to determine the relative reaches of each modulation. Even in small cells supporting 64 QAM 3/4 CTC to the cell edge, WiMAX deployments will only support 20 households per sector when tested against 2015 requirements. Range-limited cells in rural deployments, requiring QPSK at the cell edge, will support as few as 5 households per sector based on 2015 requirements.

7 References

[1] FCC Form 477

[2] Adtran, "Defining Broadband Speeds: An Analysis of Peak vs. Sustained Data Rates in Network Access Architectures," April 2009.http://insidedigitalmedia.com/internet-video-viewership-statistics/

[3] http://insidedigitalmedia.com/internet-video-viewership-statistics/

[4] Bardzell, J., Bardzell, S., and Pace, T., "Emotion, Engagement, and Internet Video," December 2008, available at http://www.scribd.com/doc/8948430/emotion-engagement-and-internet-video

[5] Limaye, P., Glapa, M., El-Sayed, M., and Gagen, P., "Impact of Bandwidth Demand growth on HFC Network," Networks 2008 presentation, available at http://www.networks2008.org/data/upload/file/Technical/A2_3_Limaye_Glapa_El-Sayed_Gagen.pdf

[6] http://www.vudu.com/product_overview.html

[7] http://www.roku.com/default.aspx

[8] http://www.popcornhour.com

[9] Keefe, K., "Maximizing Bandwidth, Minimizing Capex with DOCSIS 3.0 and an Integrated CMTS," Converge! Network Digest, August 22, 2008

[10] Wang, F., Ghosh, A., Sankaran, C., Fleming, P., Hseih, F., and Benes, S., "Mobile WiMAX Systems: Performance and Evolution," IEEE Communications Magazine, October 2008

[11] Cisco, "Cisco Visual Networking Index – Forecast and Methodology, 2007-2012," available at http://www.cisco.com/en/US/netsol/ns827/networking_solutions_sub_solution.html

[12] Cisco, "Approaching the Zettabyte Era," available at http://www.cisco.com/en/US/netsol/ns827/networking_solutions_sub_solution.html

[13] Gerber, A., Houle, J., Nguyen, H., Roughan, M., and Sen, S., "P2P, The Gorilla in the Cable," 2003, available at http://www.research.att.com/~sen/pub/p2pCable2003.final.pdf

[14] Floyd, S. and Paxson, V., "Difficulties in Simulating the Internet," IEEE/ACM Transactions on Networking, 2001, volume 9, pp 392-403

[15] Martin, J. and Westall, J., "Assessing the Impact of BitTorrent on DOCSIS Networks," Fourth International Conference on Broadband Communications, Networks and Systems, 2007

[16] Erman, D and Popescu, A., "BitTorrent Traffic Characteristics," International Multi-Conference on Computing in the Global Information Technology, 2006

[17] Qi, J., Zhang, H., and Ji, Z., "Analyzing BitTorrent Traffic Across Large Network Cyberworlds," Proceedings of the 2008 International Conference on Cyberworlds, pp.759-764, 2008

[18] http://www.joost.com/

[19] Alhaisoni, M. and Liotta, A., "Characterization of Signaling and Traffic in Joost," Peer-to-Peer Networking and Applications, Volume 2, Number 1 / March, 2009, pp. 75-83

[20] Cheshire, S., "Latency and the Quest for Interactivity," November 1996, available at http://www.stuartcheshire.org/papers/LatencyQuest.html.

[21] Marques, H., Rocha, L., Guerra, P., Almeida, J., Meira, W., and Almeida, V., "Characterizing Broadband User Behavior," Proceedings of the 2004 ACM Workshop on Next-Generation Residential Broadband Challenges, October 15-15, 2004

[22] Fukuda, K., Cho, K. and Esaki, H., "The Impact of Residential Broadband Traffic on Japanese ISP Backbones," ACM SIGCOMM Computer Communications Review, Volume 35, Number 1, January 2005

[23] Erramilli, A., Roughan, M., Veitch, D. and Willinger, W., "Self-Similar Traffic and Network Dynamics," Proceedings of the IEEE, Vol. 90, No. 5, May 2002, pp. 800-819

[24] Yu, B. and Fei, H., "Fractal Analysis of User Sessions Inter-Transaction Time in Social Networks," 4th International Conference on Wireless Communications, Networking and Mobile Computing**,** 12-14 Oct. 2008

[25] Mori, T., Kawahara, R., Naito, S. and Goto, S., "On the Characteristics of Internet Traffic Variability: Spikes and Elephants," Proceedings of the 2004 International Symposium on Applications and the Internet

[26] http://www.dtc.umn.edu/mints/home.php

[27] http://www.census.gov/

[28] Horrigan, J., "Home Broadband Adoption 2008," Pew Internet & American Life Project, available at http://www.pewinternet.org/

[29] Etemad, K., “Overview of Mobile WiMAX Technology and Evolution,” IEEE Communications Magazine, October 2008, pp. 31-40